

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicants:	Robert A. Cochran, et al.	Examiner:	Matthew A. Bradley
Serial No.:	10/798,962	Group Art Unit:	2187
Filed:	March 12, 2004	Docket No.:	10019728-1
Title:	Data Synchronization for Two Data Mirrors with Sidefiles		

APPEAL BRIEF UNDER 37 C.F.R. § 41.37

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This Appeal Brief is filed in response to the Final Office Action mailed January 11, 2008 and Notice of Appeal filed on April 11, 2008.

AUTHORIZATION TO DEBIT ACCOUNT

It is believed that no extensions of time or fees are required, beyond those that may otherwise be provided for in documents accompanying this paper. However, in the event that additional extensions of time are necessary to allow consideration of this paper, such extensions are hereby petitioned under 37 C.F.R. § 1.136(a), and any fees required (including fees for net addition of claims) are hereby authorized to be charged to Hewlett-Packard Development Company's deposit account no. 08-2025.

I. REAL PARTY IN INTEREST

The real party in interest is Hewlett-Packard Development Company, LP, a limited partnership established under the laws of the State of Texas and having a principal place of business at 20555 S.H. 249 Houston, TX 77070, U.S.A. (hereinafter "HPDC"). HPDC is a Texas limited partnership and is a wholly-owned affiliate of Hewlett-Packard Company, a Delaware Corporation, headquartered in Palo Alto, CA. The general or managing partner of HPDC is HPQ Holdings, LLC.

II. RELATED APPEALS AND INTERFERENCES

There are no known related appeals, judicial proceedings, or interferences known to appellant, the appellant's legal representative, or assignee that will directly affect or be directly affected by or have a bearing on the Appeal Board's decision in the pending appeal.

III. STATUS OF CLAIMS

Claims 1 – 32 are pending in the application and stand finally rejected. Claim 33 was canceled. The rejection of claims 1 – 32 is appealed.

IV. STATUS OF AMENDMENTS

No amendments were made after receipt of the Final Office Action. All amendments have been entered.

V. SUMMARY OF CLAIMED SUBJECT MATTER

The following provides a concise explanation of the subject matter defined in each of the claims involved in the appeal, referring to the specification by page and line number and to the drawings by reference characters, as required by 37 C.F.R.

§ 41.37(c)(1)(v). Each element of the claims is identified by a corresponding reference to the specification and drawings where applicable. Note that the citation to passages in the specification and drawings for each claim element does not imply that the limitations from the specification and drawings should be read into the corresponding claim element or that these are the sole sources in the specification supporting the claim features.

Claim 1

A data synchronization method for a redundant data storage arrangement in which there are at least a primary storage entity and mirroring first and second remote storage entities in communication therewith, respectively, writes to the first and second remote storage entities being tracked via respective first and second sidefiles, the first and second storage entities having different levels of write-currency relative to each other, the method comprising (pages 5-6, [0018-0022]):

receiving acknowledgements from the second remote storage entity at both the primary storage entity and the first remote storage entity ([ACKs received at both the primary and first remote storage entities: p. 10, [0034-0035]);

comparing acknowledgements and sequence numbers in the first sidefile with acknowledgements and sequence numbers in the second sidefile (Figs. 2A, 2B illustrate the comparison: see p. 10 and 13 [0034-0036, 0044]); and

updating writes stored at the second remote storage entity based upon the comparison of the first and second sidefiles (the write-update is based on a comparison of sequence numbers in sidefiles 22 and 32: see p. 11, [0039]).

Claim 13

A redundant data storage arrangement comprising:

a primary storage entity (#10) to (1) forward writes to each of a mirroring first and second remote storage entity (#s 20, 30) and (2) forward acknowledgements from the

second remote storage entity to the first remote storage entity (see pages 6-7 at [0024] and p. 10 at [0035]);

the mirroring first remote storage entity, in communication with the primary storage entity, which includes a first sidefile (#22) via which writes forwarded from the primary storage entity are tracked (see p. 8 [0027] for discussion of sidefiles, and p. pages 9-10 at [0032]); and

the mirroring second remote storage entity, in communication with the primary storage entity, which includes a second sidefile (#32) via which writes forwarded from the primary storage are tracked (see p. 8 [0027] for discussion of sidefiles, and p. pages 9-10 at [0032]);

the first and second storage entities having different levels of write-currency relative to each other (see p. 6-7 discussing different levels of write-currency at [0024]);

an initiating one of the first and second remote storage entities being operable to compare acknowledgements and sequence numbers in the first and second sidefiles (Figs. 2A, 2B illustrate the comparison: see p. 10 and 13 [0034-0036, 0044]), and

invoke an updating of writes stored at the second remote storage entity based upon the comparison of the first and second sidefiles (the write-update is based on a comparison of sequence numbers in sidefiles 22 and 32: see p. 11, [0039]).

Claim 28

A data synchronization method for a redundant data storage arrangement in which there are at least a primary storage entity and mirroring first and second remote storage entities in communication therewith, respectively, the method comprising:

synchronously receiving writes at the first remote storage entity that have been forwarded from the primary storage entity (see p. 8, [0228]);

maintaining a sidefile via which are tracked items that include (see p. 8, [0027])

acknowledgements from the second remote storage entity received at both the primary storage entity and the first remote storage entity ([ACKs received at both the primary and first remote storage entities: p. 10, [0034-0035]),

sequence numbers assigned to writes (see pages 8-9, [0029]),

writes received at the first remote storage entity (writes received at storage entity, see paragraphs [0027, 0031, 0036]), and

writes received at the second remote storage entity (writes received at storage entity, see paragraphs [0027, 0033]); and

comparing both acknowledgements and sequence numbers at the first remote storage entity to determine whether to update writes stored at the second remote storage entity after a failure of the primary storage entity (Figs. 2A, 2B illustrate the comparison: see p. 10 and 13 [0034-0036, 0044]).

Claim 29

A mirroring first remote storage entity for a redundant data storage arrangement in which there are at least the first remote storage entity (#22) and a primary storage entity (#10) in communication therewith, and a mirroring second remote storage entity in communication with the primary storage entity, respectively, the first remote storage entity comprising:

memory to store data (see p. 15 at paragraph [0050]); and

a processor operable to (see p. 15 at paragraph [0050])

synchronously receive writes that have been forwarded from the primary storage entity (see pages 6-7 at paragraph [0024] discussing synchronously receiving writes);

maintain a sidefile in the memory via which are tracked items that include acknowledgements from the second remote storage entity received at both the primary storage entity and the first remote storage entity ([ACKs received at both the primary and first remote storage entities: p. 10, [0034-0035]),

sequence numbers assigned to writes (see pages 8-9 at paragraph [0029]),

writes received at the first remote storage entity (see paragraphs [0027, 0031, 0036]), and

writes received at the second remote storage entity (see paragraphs [0027, 0033]); and

compare both acknowledgements and sequence numbers at the first remote storage entity to determine whether to update writes stored at the second remote storage entity after a failure of the primary storage entity (Figs. 2A, 2B illustrate the comparison: see p. 10 and 13 [0034-0036, 0044]).

Claim 30

A data synchronization method for a redundant data storage arrangement in which there a primary storage entity and mirroring first and second remote storage entities in communication therewith, respectively, the method comprising:

synchronously forwarding to the first remote storage entity writes from the primary storage entity (see pages 6-7 at [0024] for synchronously forwarding writes);
informing the first remote storage entity regarding writes acknowledged to have been received at the second remote storage entity (see p. 10 at paragraphs [0034-0035] for informing when writes acknowledged, ACK); and

comparing (1) the writes acknowledged to have been received and (2) sequence numbers assigned to the writes to determine data content at the second remote storage entity (Figs. 2A, 2B illustrate the comparison: see p. 10 and 13 [0034-0036, 0044]).

Claim 31

A primary storage entity for a redundant data storage arrangement in which there are at least the primary storage entity and mirroring first and second remote storage entities in communication therewith, respectively, the primary storage entity comprising:

memory to store data (see p. 15 at paragraph [0050]); and

a processor operable to (see p. 15 at paragraph [0050])

synchronously forward writes to the first remote storage entity (see pages 6-7 at paragraph [0024] discussing synchronously forwarding writes);

receive indications of writes acknowledged to have been received at the second remote storage entity (see pages 6-7 at paragraph [0024] discussing synchronously receiving writes);

inform the first remote storage entity regarding the indications ([ACKs received at both the primary and first remote storage entities: (see p. 10, [0034-0035]);

assign sequence numbers to writes (see pages 8-9 at paragraph [0029]);
send the sequence numbers to the first and second remote storage entities
(see p. 9 at paragraph [0031]); and

compare both acknowledgements and sequence numbers at the first
remote storage entity to determine whether to update writes stored at the second
remote storage entity after a failure of the primary storage entity (Figs. 2A, 2B
illustrate the comparison; see p. 10 and 13 [0034-0036, 0044]).

Claim 32

A redundant data storage arrangement comprising:

primary storage means (example means is #10) for storing writes received from a
host (#5, see pages 6-7 at paragraph [0024]);

first remote mirror means (example means is #20) for mirroring writes forwarded
from the primary storage means and for tracking such writes via a first sidefile (#22, see
page 8 at [0027] and p. 9 at [0032]); and

second remote mirror means (example means is #30) for mirroring writes
forwarded from the primary storage means and for tracking writes such writes via a
second sidefile (#32, see page 8 at [0027] and p. 9 at [0032]);

the first and second storage entities having different levels of write-currency
relative to each other (see paragraphs [0027, 0032]);

the first sidefile including (1) sequence numbers assigned to writes and (2)
acknowledgements from the second remote mirror means (see pages 8-9 at paragraph
[0029] and p. 9 at paragraph [0031]);

the second sidefile including sequence numbers assigned to writes (see pages 8-9
at paragraph [0029] and p. 9 at paragraph [0031]);

comparison means, responsive to the primary storage means being rendered
inoperative, for comparing the first and second sidefiles (example means is first remote
storage #20: Figs. 2A, 2B illustrate the comparison; see p. 10 and 13 [0034-0036, 0044]),
and

update means, responsive to the comparison, for updating writes stored at the
second remote mirror means based upon the comparison of the first and second sidefiles

(example means is first remote storage #20: the write-update is based on a comparison of sequence numbers in sidefiles 22 and 32: see p. 11, [0039]).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

Claims 1-4, 6, 13-17, 19, and 26-32 are rejected under 35 USC § 103(a) as being unpatentable over USPN 2005/0038968 (Iwamura) in view of USPN 2002/0083281 (Carteau).

Claims 5 and 18 are rejected under 35 USC § 103(a) as being unpatentable over Iwamura in view of Carteau and USPN 6,912,483 (Frederick).

Claims 7-9 and 20-22 are rejected under 35 USC § 103(a) as being unpatentable over Iwamura in view of Carteau and USPN 6,260,125 (McDowell).

Claims 10-12 and 23-25 are rejected under 35 USC § 103(a) as being unpatentable over Iwamura in view of Carteau and USPN 6,098,179 (Harter).

VII. ARGUMENT

The rejection of claims 1 – 32 is improper, and Applicants respectfully request reversal of these rejections.

The claims do not stand or fall together. Instead, Applicants present separate arguments for various claims. Each of these arguments is separately argued below and presented with separate headings and sub-heading as required by 37 C.F.R.

§ 41.37(c)(1)(vii).

Overview of Claims and Primary References (Iwamura and Carteau)

As a precursor to the arguments, Appellants provide an overview of the claims and the primary references (Iwamura and Carteau). This overview will assist in determining the scope and content of the prior art as required in *Graham* (see *Graham v. John Deere Co. of Kansas City*, 383 U.S. 1, 17-18 setting out an objective analysis for applying 103 rejections).

The claims are generally directed to redundant data storage that includes a primary storage entity and mirroring first and second remote storage entities having different levels of write-currency relative to each other. Writes to the first and second remote storage entities are tracked through sidefiles. The second remote storage entity sends acknowledgements to both the primary storage entity and the first remote storage. Further, acknowledgements and sequence numbers are compared in the first sidefile with acknowledgements and sequence numbers in the second sidefile. Writes stored at the second remote storage entity are updated based upon the comparison of the first and second sidefiles.

Iwamura uses logs stored in a log storage area to perform a rollback for database recovery ([0257]). The logs and rollback are based on time: “Accordingly, when the database is recovered, the host 110 starts the recovery processing from the log at this time” ([0235]). Iwamura does mention an alternative embodiment wherein a sequential ID is used: “Here, the case of the ID has a feature that if a rule is set such that continuous numbers are previously assigned, a missing number can be recognized” ([0236]).

Iwamura does not teach or suggest using acknowledgements as part of his rollback process for database recovery.

Carteau teaches a computer coupled to a disk subsystem having two controllers. The controllers are coupled to plural disk drives. After a controller receives a write request from a host, the controller writes the request to a disk and transmits the request to the second controller. The second controller sends an acknowledgement to the first controller which then acknowledges the write to the host. **Carteau does not teach or suggest using these acknowledgements as part of a process for database recovery.**

Claim Rejections: 35 USC § 103(a)

Claims 1-4, 6, 13-17, 19, and 26-32 are rejected under 35 USC § 103(a) as being unpatentable over USPN 2005/0038968 (Iwamura) in view of USPN 2002/0083281 (Carteau). Applicants respectfully traverse these rejections.

The claims recite one or more elements that are not taught or suggested in Iwamura in view of Carteau. These missing elements show that the differences between the combined teachings in the art and the recitations in the claims are great. As such, the pending claims are not a predictable variation of the art to one of ordinary skill in the art.

Claim 1

As one example, claim 1 recites “comparing acknowledgements and sequence numbers in the first sidefile with acknowledgements and sequence numbers in the second sidefile.” In other words, **claim 1 recites comparing both acknowledgments and sequence numbers in one sidefile with acknowledgements and sequence numbers in a second sidefile.** This comparison of both acknowledgements and sequence numbers is used to update writes stored at the second remote storage entity.

As noted above, Iwamura uses logs based on time and mentions the use of continuous numbers as an alternative. Iwamura, though, never teaches or suggests comparing both acknowledgements and sequence numbers in one sidefile with acknowledgements and sequence numbers in a second sidefile as recited in claim 1. Carteau discusses transmitting acknowledges, but never teaches or suggests comparing

both acknowledgements and sequence numbers in one sidefile with acknowledgements and sequence numbers in a second sidefile as recited in claim 1.

The differences between the claims and the teachings in the art are great since the references fail to teach or suggest all of the claim elements. As such, the pending claims are not a predictable variation of the art to one of ordinary skill in the art.

For at least these reasons, claim 1 and its dependent claims are allowable over Iwamura.

Claim 13

As one example, claim 13 recites comparing both “acknowledgements and sequence numbers in the first and second sidefiles.” As noted above, Iwamura uses logs based on time and mentions the use of continuous numbers as an alternative. Iwamura, though, never teaches or suggests comparing both acknowledgements and sequence numbers as recited in claim 13. Carteau discusses transmitting acknowledges, but never teaches or suggests comparing both acknowledgements and sequence numbers as recited in claim 13.

The differences between the claims and the teachings in the art are great since the references fail to teach or suggest all of the claim elements. As such, the pending claims are not a predictable variation of the art to one of ordinary skill in the art.

For at least these reasons, claim 13 and its dependent claims are allowable over Iwamura.

Claims 28 and 29

By way of example, claims 28 and 29 recite “comparing both acknowledgements and sequence numbers at the first remote storage entity to determine whether to update writes stored at the second remote storage entity after a failure of the primary storage entity.” As noted above, Iwamura uses logs based on time and mentions the use of continuous numbers as an alternative. Iwamura, though, never teaches or suggests comparing both acknowledgements and sequence numbers as recited in claims 28 and 29. Carteau discusses transmitting acknowledges, but never teaches or suggests comparing both acknowledgements and sequence numbers as recited in claims 28 and 29.

The differences between the claims and the teachings in the art are great since the references fail to teach or suggest all of the claim elements. As such, the pending claims are not a predictable variation of the art to one of ordinary skill in the art.

For at least these reasons, claims 28 and 29 and their dependent claims are allowable over Iwamura.

Claim 30

By way of example, claim 30 recites “comparing (1) the writes acknowledged to have been received and (2) sequence numbers assigned to the writes to determine data content at the second remote storage entity.” As noted above, Iwamura uses logs based on time and mentions the use of continuous numbers as an alternative. Iwamura, though, never teaches or suggests comparing both acknowledgements and sequence numbers as recited in claim 30. Carteau discusses transmitting acknowledgements but never teaches or suggests comparing both acknowledgements and sequence numbers as recited in claim 30.

The differences between the claims and the teachings in the art are great since the references fail to teach or suggest all of the claim elements. As such, the pending claims are not a predictable variation of the art to one of ordinary skill in the art.

For at least these reasons, claim 30 is allowable over Iwamura.

Claim 31

By way of example, claim 31 recites a processor that is operable to “compare both acknowledgements and sequence numbers at the first remote storage entity to determine whether to update writes stored at the second remote storage entity after a failure of the primary storage entity.” As noted above, Iwamura uses logs based on time and mentions the use of continuous numbers as an alternative. Iwamura, though, never teaches or suggests comparing both acknowledgements and sequence numbers as recited in claim 31. Carteau discusses transmitting acknowledges, but never teaches or suggests comparing both acknowledgements and sequence numbers as recited in claim 31.

The differences between the claims and the teachings in the art are great since the references fail to teach or suggest all of the claim elements. As such, the pending claims are not a predictable variation of the art to one of ordinary skill in the art.

For at least these reasons, claim 31 is allowable over Iwamura.

Claim 32

As one example, claim 32 recites that the first sidefile includes both sequence numbers and acknowledgements. The second sidefile also includes sequence numbers. A comparison means compares the first and second sidefiles. As noted above, Iwamura uses logs based on time and mentions the use of continuous numbers as an alternative. Iwamura, though, never teaches or suggests comparing both acknowledgements and sequence numbers as recited in claim 32. Carteau mentions acknowledgements but does not teach or suggest comparing both acknowledgements and sequence numbers as recited in claim 32.

The differences between the claims and the teachings in the art are great since the references fail to teach or suggest all of the claim elements. As such, the pending claims are not a predictable variation of the art to one of ordinary skill in the art.

For at least these reasons, claim 32 is allowable over Iwamura.

Claim Rejections: 35 USC § 103(a)

Claims 5 and 18 are rejected under 35 USC § 103(a) as being unpatentable over Iwamura in view of Carteau and USPN 6,912,483 (Frederick). These rejections are traversed.

As noted above, Iwamura and Carteau do not teach or suggest all the elements of independent claims 1 and 13. Frederick does not cure the deficiencies of Iwamura and Carteau. For at least these reasons, claims 5 and 18 are allowable over Iwamura, Carteau, and Frederick.

Claim Rejections: 35 USC § 103(a)

Claims 7-9 and 20-22 are rejected under 35 USC § 103(a) as being unpatentable over Iwamura in view of Carteau and USPN 6,260,125 (McDowell). These rejections are traversed.

As noted above, Iwamura and Carteau do not teach or suggest all the elements of independent claims 1 and 13. McDowell does not cure the deficiencies of Iwamura and Carteau. For at least these reasons, claims 7-9 and 20-22 are allowable over Iwamura, Carteau, and McDowell.

Claim Rejections: 35 USC § 103(a)

Claims 10-12 and 23-25 are rejected under 35 USC § 103(a) as being unpatentable over Iwamura in view of Carteau and USPN 6,098,179 (Harter). These rejections are traversed.

As noted above, Iwamura and Carteau do not teach or suggest all the elements of independent claims 1 and 13. Harter does not cure the deficiencies of Iwamura and Carteau. For at least these reasons, claims 10-12 and 23-25 are allowable over Iwamura, Carteau, and Harter.

CONCLUSION

In view of the above, Applicants respectfully request the Board of Appeals to reverse the Examiner's rejection of all pending claims.

Any inquiry regarding this Amendment and Response should be directed to Philip S. Lyren at Telephone No. 832-236-5529. In addition, all correspondence should continue to be directed to the following address:

Hewlett-Packard Company
Intellectual Property Administration
P.O. Box 272400
Fort Collins, Colorado 80527-2400

Respectfully submitted,

/Philip S. Lyren #40,709/

Philip S. Lyren
Reg. No. 40,709
Ph: 832-236-5529

VIII. Claims Appendix

1. A data synchronization method for a redundant data storage arrangement in which there are at least a primary storage entity and mirroring first and second remote storage entities in communication therewith, respectively, writes to the first and second remote storage entities being tracked via respective first and second sidefiles, the first and second storage entities having different levels of write-currency relative to each other, the method comprising:

receiving acknowledgements from the second remote storage entity at both the primary storage entity and the first remote storage entity;

comparing acknowledgements and sequence numbers in the first sidefile with acknowledgements and sequence numbers in the second sidefile; and

updating writes stored at the second remote storage entity based upon the comparison of the first and second sidefiles.

2. The method of claim 1, wherein the updating of writes includes forwarding to the second remote storage entity writes that are present in the first sidefile but not yet present in the second sidefile.

3. The method of claim 1, wherein the comparing is done on the basis of sequence numbers associated with the writes.

4 The method of claim 1, further comprising:

establishing a communication connection between the first and second remote storage entities in response to the primary storage entity becoming inoperative;

wherein the comparing and updating are also performed in response to the primary storage entity becoming inoperative.

5. The method of claim 1, further comprising:

configuring the first and second sidefiles to store a fixed number of writes therein, respectively; and

adding a new write to the first and second sidefiles by overwriting the oldest write therein, respectively.

6. The method of claim 1, further comprising:

adaptively adjusting the number of writes that can be stored in the first sidefile.

7. The method of claim 6, wherein the adaptive adjustment is based upon the writes that are stored in the second sidefile.

8. The method of claim 7, wherein the adaptive adjustment more particularly is based upon the write stored in the second sidefile which has the oldest contiguous sequence number.

9. The method of claim 7, further comprising:

identifying at least some of the writes stored in the second sidefile; and then

accordingly informing the first remote storage entity regarding such identities.

10. The method of claim 1, further comprising:

configuring the first sidefile to include a field that is used to track whether a write has been acknowledged by the second remote storage entity as having been received.

11. The method of claim 1, further comprising:

providing a third sidefile via which writes received thereby are tracked; and
configuring the third sidefile to include a field that is used to track whether a write has been acknowledged by the second remote storage entity as having been received.

12. The method of claim 11, further comprising:

configuring the third sidefile further to include a field that is used to track whether a write-acknowledgement forwarded from the second remote storage entity has been acknowledged as having been received by the first remote storage entity.

13. A redundant data storage arrangement comprising:

a primary storage entity to (1) forward writes to each of a mirroring first and second remote storage entity and (2) forward acknowledgements from the second remote storage entity to the first remote storage entity;

the mirroring first remote storage entity, in communication with the primary storage entity, which includes a first sidefile via which writes forwarded from the primary storage entity are tracked; and

the mirroring second remote storage entity, in communication with the primary storage entity, which includes a second sidefile via which writes forwarded from the primary storage are tracked;

the first and second storage entities having different levels of write-currency relative to each other;

an initiating one of the first and second remote storage entities being operable to compare acknowledgements and sequence numbers in the first and second sidefiles, and

invoke an updating of writes stored at the second remote storage entity based upon the comparison of the first and second sidefiles.

14. The redundant data storage arrangement of claim 13, wherein the updating is performed by the first remote storage entity, which is operable to do so by forwarding to the second remote storage entity writes that are present in the first sidefile but not yet present in the second sidefile.

15. The redundant data storage arrangement of claim 13, wherein each of the first and second remote storage entities is operable to

preserve in the respective sidefile sequence numbers associated with the writes;
and

sort the respective sidefile according to the sequence numbers.

16. The redundant data storage arrangement of claim 13, wherein:

the initiating one is further operable to establish a communication connection between itself and the other remote storage entity in response to the primary storage entity becoming inoperative; and

the comparison and the update are performed in response to the primary storage entity becoming inoperative.

17. The redundant data storage arrangement of claim 13, wherein the first remote storage entity is closer in proximity to the primary storage entity than the second remote storage entity.

18. The redundant data storage arrangement of claim 13, wherein:

the first and second sidefiles are configured to store a fixed number of writes therein, respectively; and

the first and second remote storage entities are operable to add a new write to the first and second sidefiles by overwriting the oldest write therein, respectively.

19. The redundant data storage arrangement of claim 13, wherein:

the first remote storage entity is operable to adaptively adjust the number of writes that can be stored in the first sidefile.

20. The redundant data storage arrangement of claim 19, wherein the adaptive adjustment is based upon the writes are stored in the second sidefile.

21. The redundant data storage arrangement of claim 20, wherein the adaptive adjustment more particularly is based upon the write stored in the second sidefile which has the newest contiguous sequence number.

22. The redundant data storage arrangement of claim 20, wherein:

the primary storage entity is operable to identify at least some of the writes stored in the second sidefile and then accordingly inform the first remote storage entity regarding such identities.

23. The redundant data storage arrangement of claim 13, wherein:

the first sidefile includes a field that is used to track whether a write has been acknowledged by the second remote storage entity as having been received.

24. The redundant data storage arrangement of claim 13, wherein:

the primary storage entity includes a third sidefile to track writes received thereby;

the first sidefile including a field that is used to track whether a write has been acknowledged by the second remote storage entity as having been received.

25. The redundant data storage arrangement of claim 24, wherein the first sidefile further includes a field that is used to track whether a write-acknowledgement forwarded from the second remote storage entity has been acknowledged as having been received by the first remote storage entity.

26. The redundant data storage arrangement of claim 13, wherein each of the first and second remote storage entities represents a tracked write in the first and second sidefile, respectively, with: location information as to where on a physical medium the write is to be performed; actual data associated with the write that is to be written to the physical medium; and a sequence number uniquely associated with the write.

27. The redundant data storage arrangement of claim 13, wherein:

the first remote storage entity receives writes forwarded synchronously from the primary storage entity; and

the second remote storage entity receives writes forwarded asynchronously from the primary storage entity.

28. A data synchronization method for a redundant data storage arrangement in which there are at least a primary storage entity and mirroring first and second remote storage entities in communication therewith, respectively, the method comprising:

synchronously receiving writes at the first remote storage entity that have been forwarded from the primary storage entity;

maintaining a sidefile via which are tracked items that include

acknowledgements from the second remote storage entity received at both the primary storage entity and the first remote storage entity,

sequence numbers assigned to writes,

writes received at the first remote storage entity, and

writes received at the second remote storage entity; and
comparing both acknowledgements and sequence numbers at the first remote storage entity to determine whether to update writes stored at the second remote storage entity after a failure of the primary storage entity.

29. A mirroring first remote storage entity for a redundant data storage arrangement in which there are at least the first remote storage entity and a primary storage entity in communication therewith, and a mirroring second remote storage entity in communication with the primary storage entity, respectively, the first remote storage entity comprising:

memory to store data; and
a processor operable to
synchronously receive writes that have been forwarded from the primary storage entity;
maintain a sidefile in the memory via which are tracked items that include
acknowledgements from the second remote storage entity received
at both the primary storage entity and the first remote storage entity,
sequence numbers assigned to writes,
writes received at the first remote storage entity, and
writes received at the second remote storage entity; and
compare both acknowledgements and sequence numbers at the first remote storage entity to determine whether to update writes stored at the second remote storage entity after a failure of the primary storage entity.

30. A data synchronization method for a redundant data storage arrangement in which there a primary storage entity and mirroring first and second remote storage entities in communication therewith, respectively, the method comprising:

synchronously forwarding to the first remote storage entity writes from the primary storage entity;

informing the first remote storage entity regarding writes acknowledged to have been received at the second remote storage entity; and

comparing (1) the writes acknowledged to have been received and (2) sequence numbers assigned to the writes to determine data content at the second remote storage entity.

31. A primary storage entity for a redundant data storage arrangement in which there are at least the primary storage entity and mirroring first and second remote storage entities in communication therewith, respectively, the primary storage entity comprising:

memory to store data; and

a processor operable to

synchronously forward writes to the first remote storage entity;

receive indications of writes acknowledged to have been received at the second remote storage entity;

inform the first remote storage entity regarding the indications;

assign sequence numbers to writes;

send the sequence numbers to the first and second remote storage entities;
and

compare both acknowledgements and sequence numbers at the first
remote storage entity to determine whether to update writes stored at the second
remote storage entity after a failure of the primary storage entity.

32. A redundant data storage arrangement comprising:

primary storage means for storing writes received from a host;

first remote mirror means for mirroring writes forwarded from the primary
storage means and for tracking such writes via a first sidefile; and

second remote mirror means for mirroring writes forwarded from the primary
storage means and for tracking writes such writes via a second sidefile;

the first and second storage entities having different levels of write-currency
relative to each other;

the first sidefile including (1) sequence numbers assigned to writes and (2)
acknowledgements from the second remote mirror means;

the second sidefile including sequence numbers assigned to writes;

comparison means, responsive to the primary storage means being rendered
inoperative, for comparing the first and second sidefiles, and

update means, responsive to the comparison, for updating writes stored at the
second remote mirror means based upon the comparison of the first and second sidefiles.

33. (canceled)

IX. EVIDENCE APPENDIX

None.

X. RELATED PROCEEDINGS APPENDIX

None.